

Health Maintenance System: An Application of Recovery Oriented Computing for HPEC Systems

Gerry Pocock

SKY Computers, Inc.

Phone: (978) 250-1920

Fax: (978) 250-0036

Email: pocock@skycomputers.com

Until recently, the primary, single aspect of HPEC systems that has been most critical has been "performance", in terms of processor speeds and I/O throughput. As processor speeds and I/O throughput has continued to increase, and as the capability to build larger and larger systems has improved, the need for raw performance is becoming less critical. Now, it is the ability to achieve a high level of application availability that is becoming as critical as performance.

In this paper, we will present a CORBA based framework upon which highly available applications can be constructed. This framework, known as the Health Maintenance System, provides the application, system managers, and management tools with the ability to "manage" all resources within a system such that the "health" of the system can be maintained. The management of these resources involves the ability to "sense" the state of the resource, to control the resource, and to run tests on the resource in order to pro-actively detect any latent problems.



The primary facet of the framework is the "resource manager". The resource managers provide local management support for all system resources. In addition, the resource managers provide management access to clients, e.g., the application. This access is provided via a set of "client interface" modules that provide a wide variety of interfaces, e.g., APIs, agents, etc. It is this combination of resource managers and client interface modules that allow the framework to be easily configured for a specific HPEC system.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 20 AUG 2004		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Health Maintenance System: An Application of Recovery Oriented Computing for HPEC Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SKY Computers				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM001694, HPEC-6-Vol 1 ESC-TR-2003-081; High Performance Embedded Computing (HPEC) Workshop(7th). , The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Health Management System: An Application of Recovery Oriented Computing (ROC) Targeted at HPEC Systems

HA Approach: Redundancy

- Goal: Increase Mean Time To Failure
- Classes
 - Dual Redundancy (Hot Fail Over)
 - Triple Redundancy (Result Comparison)
- Redundancy at System/Component Level
- Drawbacks:
 - High Costs
 - Low Density
 - Additional Complexity



SKY's HAA Approach: Recovery Oriented Computing

Two Basic Tenets:

- Failure Rates of Both Software and Hardware are Non-Negligible and Increasing
- Systems Cannot be Completely Modeled for Reliability Analysis (thus their failure modes cannot be predicted in advance)

Goal:

- Decrease Mean Time to Repair

ROC Mechanisms:

- Detection (Sensing and Diagnostics)
- Isolation
- Use of Excess Capacity (if available)
- Repair/Recovery

Applying ROC to HPEC

Hardware:

- Quality Components
- Builtin Sensing of all Major Components
- Control of all Major Components (reset, etc.)
- Excess Capacity (where possible)

OS Middleware:

- Quality Components
- Builtin Sensing of all Major Components
- Control of all Major Components

Application:

- Quality Components
- Builtin Sensing of all Major Components
- Control of all Major Components
- Overall System Management (Sensing and Control)

Hardware Support

- **HAA Support Blade**


- ✧ Tini Management Processor (Java Processor)
- ✧ I2C Integration
- ✧ TCP/IP External Access

- **Compute/IO Blades**

- ✧ Out-of-band Management Controller
- ✧ Temperature Monitoring
- ✧ Voltage Monitoring
- ✧ Heart Beat Monitor
- ✧ Power Control/Reset
- ✧ I2C Integration

- **Chassis**

- ✧ Fan Monitoring
- ✧ Voltage Monitoring
- ✧ Power Control/Reset



Health Management System (HMS)

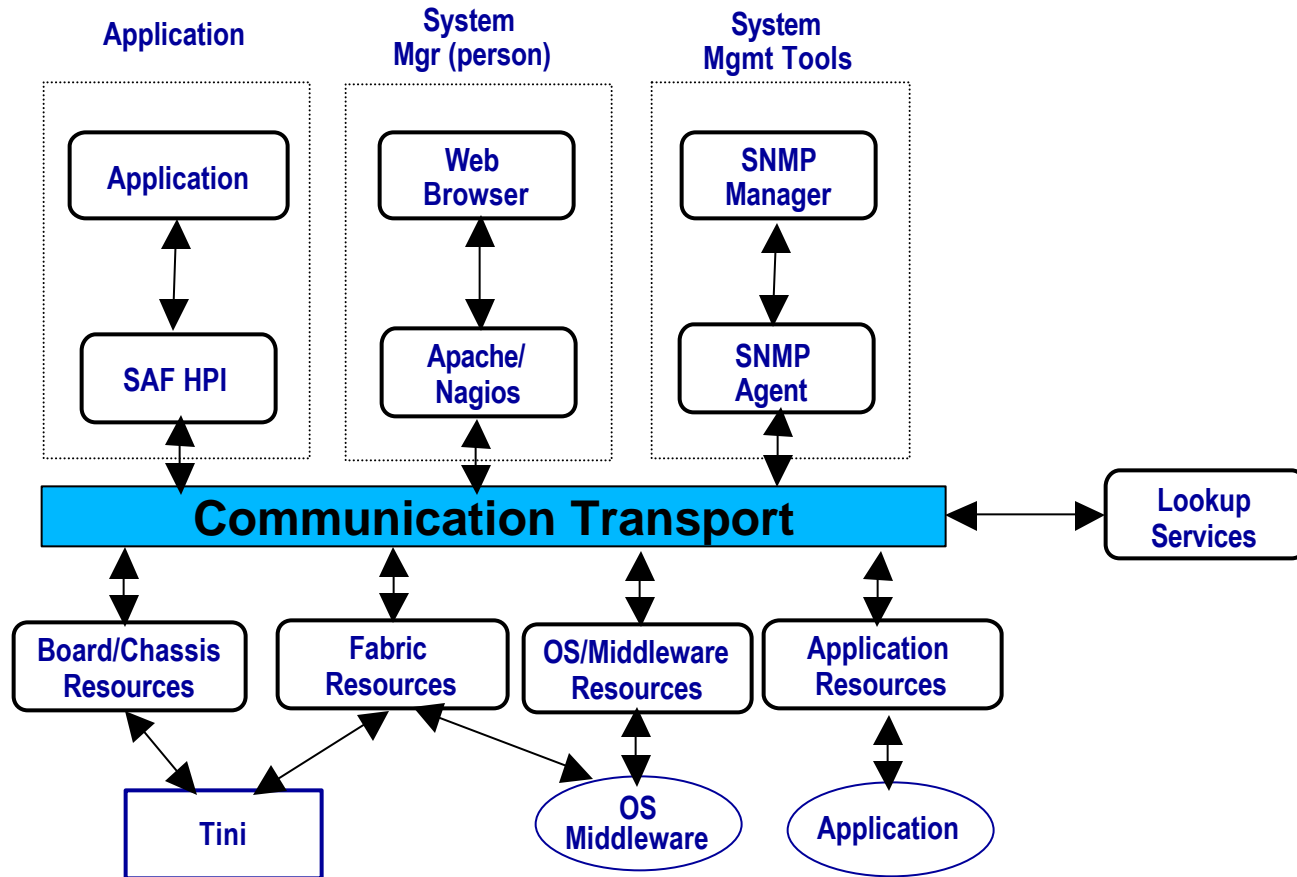
GOALS:

- Provide Capability to Instrument OS, Middleware, and Application (analogous to hardware instrumentation)
- Provide Uniform View of Entire System (hardware, OS, middleware, and application)
- Provide Integrated Diagnostics
- Provide Access Using Standard Interfaces
- Minimal Performance Impact
- Easily Extensible and Configurable (in order to meet individual application requirements)

Health Management System (HMS)

- Server Objects: Sensors, Controllers, and Timers
 - Embedded within the hardware, OS, middleware, and application
 - Combined into a Resource Object
- Clients: Application, Management Tools, and Users
- Communication: Event Driven, Request Driven, and Timer Driven Messaging
- Lookup Services
- Extensible
 - ✍ Can support an arbitrary number of servers and clients
 - ✍ Application developers can add application specific servers
- Configurable
 - ✍ Which servers and clients are to run
 - ✍ When and where they are to run

Example HMS Based System



Health Management System

- Used to Monitor Resource Usage (Development and Runtime)
 - ✧ Hardware (temperature, voltage, etc)
 - ✧ OS/Middleware (processor load, data throughput, etc)
 - ✧ Application (queue lengths, wait times, etc)
- Used to Manage These Resources
- Used to Detect and Isolate Faults
- Used to Predict Possible Future Faults
- Used to Gather Statistics on Resource Usage and Performance
- Used to Determine the Health of Resources (Diagnostics)

Future Directions

- Tight Integration with SKY Analysis Tools
- Tight Integration with SKY Development Tools
- Pattern-based Application Recovery Libraries
- Dynamic Insertion of Sensors/Controllers (Dynamic Probes)
- Support for Other Hardware Environments (Hot-Swap)